

CODE IMAGE DISTRIBUTION IN A MULTI-NODE NETWORK OF PROCESSORS

FIELD OF THE INVENTION

This invention relates to multi-node networks of embedded systems employing processors, each processor operating in accordance with a code image, and, more particularly, to insuring that each processor of the network has its code image, for example, when the processor or the network is powered on.

BACKGROUND OF THE INVENTION

Typically, each processor of a multi-node network of embedded systems employing processors has a high speed RAM (Random Access Memory) for storing a code image which, when executed, operates the processor. An example of an embedded system having a plurality of modules with processors at nodes of the system, comprises an automated data storage library, which stores removable data storage media in storage shelves, which has a plurality of data storage drives to read and/or write data on the data storage media, and which has at least one robot to transport the data storage media between the storage shelves and the data storage drives, with processors at the modules to

DOCKET: TUC920000080-US1

operate the library, and with a network to couple the modules of the embedded library system. Such a library system may comprise a hundred or more data storage drive processors and tens of library processors. It can become a real challenge to maintain
5 consistent code levels in a network of a large number of nodes.

Some multi-node networks, such as LANs (Local Area Networks), are employed for coupling together a number of like components, such as PCs (Personal Computers), with peripheral devices, such as printers. In at least one instance, such as
10 discussed by U.S. Patent No. 5,815,722, executable code files for the specific devices may be updated by downloading directly to RAM. In the patent, a communication program operates to broadcast an inquiry on the LAN to a specific network board, to receive location information of the designated board, e.g., of a
15 printer. The executable file is then directly downloaded into RAM on the designated board through the LAN. However, to prevent loss of the executable code when any of the devices or boards is powered off, the RAM is non-volatile (called NVRAM).

The executable code for printers and similar devices is
20 exceptionally small, requiring only a very small NVRAM). In more major systems, such as automated data storage libraries, a module processor may comprise a processor of the power of a workstation or PC, and the executable code for each module processor is quite large.

DOCKET: TUC920000080-US1

As a result, typically, each processor of a multi-node network has a PROM (Programmable Read-Only Memory) or a ROM (Read-Only Memory), which stores a power-on sequence, and which must store a copy of the code image to prevent loss. A power-on
5 sequence for any of the modules of an embedded system involves the execution of power-on code that is stored in the PROM or ROM device, where a module may comprise an assembly, subassembly or a circuit board. The code may first test the components of the module, and then may transfer the code image to the faster RAM.

10 Thus, the need for both NVRAMs, PROMs or ROMs and a fast RAM is very expensive. The requirement that a PROM or ROM have sufficient capacity to store the code image for the associated RAM, or that a NVRAM be used, adds significantly to the expense of each module. When multiplied by a high number of modules, the
15 expense can become a significant factor in the cost of the full system. Further, should a code image be updated, failures of the update process are not uncommon, for example, if power is lost during a code update, and may even cripple the node entirely.

SUMMARY OF THE INVENTION

An object of the present invention is to eliminate the need for a large NVRAM, PROM or ROM.

Another object of the present invention is to provide code
5 images to processors of a multi-node processing network, for example, when the processor or the network is powered on, without requiring a high capacity NVRAM, PROM or ROM.

Disclosed are a multi-node processing network, method, and computer program product, with a plurality of processors coupled
10 in the network. The processors have a minimally operational state, e.g., upon being rebooted, and have a fully operational state employing a code image.

In accordance with the present invention, the processors, when in the minimally operational state, request the code image
15 from the network. The processors may have a non-volatile memory storing code for the minimally operational state which is sufficient to at least provide the code image request. A master source is coupled in the network, the master source having at least a code image for broadcasting on the network. The master
20 source, upon receiving a code image request, waits a predetermined time period, the predetermined time period allowing any additional processor to reach the minimally operational

DOCKET: TUC920000080-US1

state. Upon completion of the predetermined time period, the master source broadcasts the code image on the network.

In another aspect of the present invention, the processors, additionally, upon the broadcast of the code image, receive and
5 implement the code image only if the processor is in the minimally operational state.

Thus, without requiring a high capacity NVRAM, PROM or ROM at each processor, the present invention provides code images to processors of a multi-node processing network, but does not
10 interrupt a currently operating processor.

A code image for the processors may be updated by first updating the master source, and rebooting all of the processors that are on the network and, upon being rebooted, the processors will request the code image, and the master source will supply
15 the updated code image. In another aspect of the present invention, in which the processors are provided for modules of a redundant system having at least two sets of redundant modules, at least one set of redundant modules is rebooted, such that the processors of the modules reach the minimally operational state
20 and request the code image from the network. The redundant modules may comprise all or part of the nodes in the complete system, and may comprise as little as a single node which is duplicated to be redundant. The master source, upon receiving a code image request, waits a predetermined time period, the

DOCKET: TUC920000080-US1

predetermined time period allowing any additional processor to reach the minimally operational state. The master source, upon completion of the predetermined time period, broadcasts the code image on the network, such that the processors requesting the

5 code image become fully operational. Then, a remaining set of redundant modules is rebooted, so that the processors of the modules reach the minimally operational state and request the code image from the network. The master source, upon receiving a code image request, waits a predetermined time period, the

10 predetermined time period allowing any additional processor to reach the minimally operational state, and, upon completion of the predetermined time period, broadcasts the code image on the network, such that the processors of the remaining set of redundant modules requesting the code image become fully

15 operational. In this manner, the code images for the entire redundant system have been updated without requiring that the operation of the entire system be interrupted.

As the result, the present invention eliminates the need to maintain and update code levels on individual processors.

20 For a fuller understanding of the present invention, reference should be made to the following detailed description taken in conjunction with the accompanying drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a block diagram of a multi-node processing network arranged in accordance with the present invention;

FIG. 2 is a diagrammatic representation of the code stored
5 at a master source of the network of FIG. 1;

FIG. 3 is an isometric view of an automated data storage library employing the network of FIG. 1;

FIG. 4 is a block diagram of the automated data storage library of FIG. 3;

10 FIG. 5 is a flow chart depicting computer implemented embodiments of the method of the present invention for distributing code images in a network system; and

FIG. 6 is a flow chart depicting a computer implemented embodiment of a method for updating code images for a redundant
15 system in accordance with the present invention;

DETAILED DESCRIPTION OF THE INVENTION

This invention is described in preferred embodiments in the following description with reference to the Figures, in which like numbers represent the same or similar elements. While this
20 invention is described in terms of the best mode for achieving

DOCKET: TUC920000080-US1

this invention's objectives, it will be appreciated by those skilled in the art that variations may be accomplished in view of these teachings without deviating from the spirit or scope of the invention.

5 FIG. 1 illustrates a multi-node processing network system 100 with a plurality of processors 105 coupled in a network 106, together with a master source 108. An example of a network comprises a multi-node system of embedded processors. The embedded system comprises a plurality of modules with processors
10 at nodes of the system, and the network serves to interconnect the modules 110 of the system, where the modules have programmable processors 105 to operate the modules and their components 111, thereby operating the system. In the illustrated example, the processors 105 each comprises a processor interface
15 112 coupling the processor in the network, a non-volatile memory 113, such as a ROM, for storing code comprising at least a boot program, a processor memory 114, such as a high speed RAM, storing a code image providing a fully operational state of the processor, and a processing unit 115 coupled to the non-volatile
20 memory 113, the processor memory 114 and the processor interface 112.

In accordance with the present invention, the processors 105 have a minimally operational state, e.g., upon being rebooted,

DOCKET: TUC920000080-US1

employing the boot program code of the non-volatile memory 113, and have a fully operational state employing a code image stored in the processor memory 114. Herein, "reboot" may comprise (a) a power-on of a processor, (b) a reset of a processor, (c) a
5 software instruction or command to cause a processor to run its boot code, or (d) any other initialization process resulting in running the boot code.

The processor memory 114 is typically a high speed RAM that stores a code image which, when executed, operates the processor.
10 The code image stored in the high speed RAM is lost whenever the processor is powered off. If a processor is powered off, the boot program stored in the non-volatile memory 113 becomes operational upon a power-on reset of the processor. As discussed above, the typical non-volatile memory of the prior art
15 additionally contains the entire code image to prevent loss of the executable code when any of the modules is powered off.

Further in accordance with the present invention, the non-volatile memory 113 is of relatively small capacity, and only contains code for the minimally operational state which is
20 sufficient to at least provide a code image request. As an example, many processors may include a flash memory 113 as an off-the-shelf product, making a separate component unnecessary. Specifically, when a processor 105 is in the minimally

DOCKET: TUC920000080-US1

operational state, the processor employs the boot program to request the code image from the network 106. The master source 108 is coupled in the network, and has at least the requested code image for broadcasting on the network.

5 The master source 108 comprises a master interface 120 coupled in the network, a master processor 125 coupled to the master interface, and a non-volatile memory 126 coupled to the master processor 125 which stores at least the code image for the processors of the network. The master source may also comprise a
10 functional node of the network with a RAM, module components, etc. In accordance with the present invention, upon receiving a code image request from a processor 105 at the master interface 120, the master source 108 waits a predetermined time period, the predetermined time period allowing any additional processor 105
15 to reach the minimally operational state, and, upon completion of the predetermined time period, the master source 108 broadcasts the code image stored in the non-volatile memory 126, via the master interface 120, on the network 106. The predetermined time period may be renewed or extended, as is discussed hereinafter.

20 Thus, without requiring a high capacity NVRAM, PROM or ROM at each processor for storing the entire code image, the present invention provides code images to processors of a multi-node processing network.

In another aspect of the present invention, the processors 105, upon the broadcast of the code image by the master source, receive and implement the code image only if the processor is in the minimally operational state. In this manner, no running
5 fully operational processor is interrupted or disrupted to interject a new code image.

Referring additionally to FIG. 2, the processors 105 of a multi-node network of embedded systems 100 are associated with modules that perform various functions in the system. Thus, the
10 processors may employ correspondingly different code images. As the result, the master source 108 of the present invention may be required to store multiple code images. An example of the code stored in non-volatile memory 126 of master source 108 is depicted in FIG. 2, and comprises a code image 130 for operating
15 the master source 108 and other code images 131-133 for operating the various processors 105. A code image may comprise a computer program product usable with a programmable computer, the computer program product having computer readable program code embedded therein.

20 In the instance where the master source 108 also comprises a functional node of the network, in one alternative, code image 130 may comprise both the master source code and the functional node code. In another alternative, the code image 130 may

DOCKET: TUC920000080-US1

comprise the master source code, and one of the code images 131-133 may comprise the functional node code. A redundant master source 108 may also be provided and may have the same, a different, or no functional node code.

5 The master source may distribute the code images 131-133 either by 1) providing and broadcasting one code image for any code image request, where the one master source code image comprises a combination of the different code images 131-133, or by 2) selecting and broadcasting only the requested one of the
10 different code images. In the event of a dual master source, code image 130 may also be broadcast, or separate master source and module function code images may be provided and broadcast. For example, one master source may be mounted at a module of a first type and the dual master source at a module of a second
15 type, and another module function may be a processor 105. Where the one broadcast master source code image comprises a combination of different code images, the processors 105 additionally select and implement one of the combination of different code images, specifically, the code image that is
20 correct for the processor. Where the master source 108 selects the desired code image 131-133 in response to a specific request, and broadcasts the selected code image, each of the requesting processors 105 determines whether the broadcast code image is

DOCKET: TUC920000080-US1

correct for the requesting processor, and the requesting processor selects the broadcast code image for implementation if the determination is that the code image is correct for the processor. A dual master may also request and receive a code
5 image 130. Alternatively, all processors 105 may use one common code image, since they contain many common code elements, but each processor can sense node type and not execute some portions of the code.

Referring to FIGS. 3 and 4, an example of an embedded system
10 having a plurality of modules with processors at nodes of the system, comprises an automated data storage library 10, which stores removable data storage media 14 in storage shelves 16. An example of an automated data storage library is the IBM 3494 Tape Library Dataserver. The exemplary library comprises a base frame
15 11, may additionally comprise one or more extension frames 12, and may comprise a high availability frame 13. The base frame 11 of the library 10 comprises a plurality of data storage drives 15 to read and/or write data on the data storage media, and has a robot accessor 18 to transport the data storage media between the
20 storage shelves and the data storage drives. The robot 18 includes a gripper assembly 20 and may include a bar code scanner or reading system 22 to read identifying information about the data storage media 14. The library may also comprise an operator

DOCKET: TUC920000080-US1

panel 23 or other user interface which allows a user to interact with the library.

The extension frame 12 comprises additional storage shelves, and may comprise additional data storage drives 15. The high availability frame 13 may also comprise additional storage shelves and data storage drives 15, and comprises a second robot accessor 28, which includes a gripper assembly 30 and may include a bar code scanner 32 or other reading device, and an operator panel 280 or other user interface. In the event of a failure or other unavailability of the robot accessor 18, or its gripper 20, etc., the second robot accessor 28 may take over. Additionally, both robot accessors may operate simultaneously in different sections of the library to speed storage and retrieval of the data storage media.

15 The library 10 receives commands from one or more host systems 40, 41 or 42. The host systems, such as host servers, communicate with the library, either directly on path 80, or through one or more data storage drives 15 on paths 81 or 82. The paths 80, 81 or 82 may, for example, comprise SCSI busses or 20 fibre channel arbitrated loops.

The embedded system comprises processors at the modules to operate the library, with a network to couple the modules of the embedded library system. Such a library system may comprise a

DOCKET: TUC920000080-US1

hundred or more data storage drive processors and tens of library processors. Thus, the embedded system may comprise communication processor nodes 50, 155, and 250 to receive commands from the hosts and convert the commands to physical movements of the robot
5 accessors 18, 28. The communication processor nodes 50, 155, and 250 may also provide a communication link for operating the data storage drives 15. Work processor nodes 52 and 252 are located at the respective robot accessor 18 and 28, and respond to host commands received from the communication processor nodes and
10 direct the operation of the accessor, providing XY move commands. The work processor nodes are also coupled to the respective scanners 22, 32. XY processor nodes 55 and 255 are located at an XY system of the accessor and respond to the XY move commands, operating the XY system to position the grippers 20, 30 to access
15 the media.

Also, operator panel processor nodes 59, 259 may be provided at the respective operator panels 23, 280 for providing an interface for communicating between the operator panel and the communication processor nodes 50, 155, and 250, the work
20 processor nodes 52, 252, and the XY processor nodes 55, 255.

The network comprises base frame common bus 60, coupling the communication processor node 50, the operator panel processor node 59, the work processor node 52, and the XY processor node

DOCKET: TUC920000080-US1

55. The extension frame 12 is coupled by an extension common bus 152 to the base frame common bus 60. The communication processor node 155 is coupled to the extension common bus 152, and communicates with both data storage drives 15 and with hosts, 5 either directly or indirectly, at input 156. Additional extension frames employing identical communication processor nodes 155, storage shelves 16, data storage drives 15, and extension busses 152, may be provided, and each is coupled to the adjacent extension frame. The high availability frame 13 10 comprises an extension common bus 200 coupled to the extension common bus 152 of an extension frame or, if there is no extension frame, to the common bus 60 of the base frame. The high availability frame extension common bus 200 couples the communication processor node 250, which is coupled to the data 15 storage drives 15 and may receive commands from hosts at input 256, to the work processor node 252, to the operator panel processor node 259, and to the XY processor node 255.

Thus, each of the processor nodes is coupled to the common busses 60, 152, 200 to form a network.

20 The common busses 60, 152, 200 may comprise a wiring network, such as the commercially available "CAN" bus system, which is a multi-drop network, having a standard access protocol and wiring standards, for example, as defined by CiA, the CAN in

DOCKET: TUC920000080-US1

Automation Association, Am Weich selgarten 26, D-91058 Erlangen, Germany. Each of the extension common busses 152, 200 may comprise a flex cable connection to the preceding adjacent common bus. Other similar bus networks may be employed for implementing 5 the present invention. Alternatively, the common busses 60, 152, 200 may comprise a wireless network system, such as RF or infrared, as is known to those of skill in the art.

The processors coupled to the network may comprise a processor of the power of a workstation or PC, in the case of the 10 communication processor nodes 50, 155, 250, the work processor nodes 52, 252, the XY processor nodes 55, 255, and each of the drives 15, and the code image for each module processor may be quite large. Also, each of the operator panel processor nodes 59, 259 may comprise a module processor requiring a code image, 15 or may comprise a separate processing system. Further, each of the types of modules conduct different jobs, potentially requiring different code images.

Referring to FIG. 1, when a module 110 or its processor 105 is repaired, replaced or upgraded, it must be powered off, the 20 repair, etc. work conducted, and then it is turned on, called a power-on reset. When the processor is powered off, the code image is lost from the RAM 114 and must be replaced before the processor, and therefore the module, can become operational.

DOCKET: TUC920000080-US1

Referring additionally to FIG. 5, the present invention provides code images to processors of a multi-node processing network without requiring a high capacity NVRAM, PROM or ROM at each processor to store the code image in the event the module is
5 powered off.

In step 300, the code image, or code images 131-133, or code image 130 in the case a dual master, of FIG. 2, are loaded to the master source 108 and, in step 301, are stored in the non-volatile memory 126 of the master source. In FIG. 1-5, the
10 master source 108 may comprise any of the processors coupled to the base frame common bus 60, and may, for example, comprise the work processor node 52. The code images 131-133 each comprises the current updated code image for the corresponding processors 105. The code images are thus ready to be broadcast when
15 requested by one or more processors.

As discussed above, a processor 105 may be rebooted, depicted as step 305, such that the boot program stored in the non-volatile memory 113 becomes operational, and is provided to the processing unit 115 of the processor 105 in step 307. In
20 accordance with the present invention, the non-volatile memory 113 is of relatively small capacity, and only contains the boot code. The boot code may first conduct at least a basic system test which tests the components of the module, in step 309, and

DOCKET: TUC920000080-US1

then places the processor 105 in the minimally operational state in step 310. The processor may previously have been powered off and subsequently powered on and is, at that time, absent a code image required to become fully operational. Then, in step 311,
5 the processor employs the boot program to request the code image from the network 106.

The master source receives the first code image request in step 320. In accordance with the present invention, upon receiving a first code image request from a processor 105 at the
10 master interface 120, the master source 108, in step 321, waits a predetermined time period. The predetermined time period of step 321 is sufficient to allow any additional processor 105 that might be powered on at about the same time, to reach the minimally operational state, and provide a code image request,
15 received at step 322. The predetermined time may be fixed or variable, and may be timed from the first received request, or alternatively from the most recently received request. Thus, as an example, if several data storage drives 15 were replaced or upgraded at the same time, they could be powered on at about the
20 same time, or a frame could be powered on and the master source 108 waits for all processors to reach the minimally operational state so that the code image will only be broadcast once.

DOCKET: TUC920000080-US1

In the event the predetermined time period is timed from the most recently received request, step 324 determines whether a new time period is required. If so, the previous time period is renewed in step 321, or, alternatively, a new predetermined time period is established in step 321, in effect extending the predetermined time period. As an example, the new predetermined time period may be shorter than the previous time period by a fixed or variable amount. If no new request is received, or if the timing is only from the first received request, the process proceeds to step 323. Step 323 determines whether the time period has expired, and, if not, continues the wait step 321.

Referring additionally to FIG. 2, as discussed above, the processors 105 of a multi-node network of embedded systems 100 are associated with modules that perform various functions in the system, and may employ correspondingly different code images. As the result, the master source 108 of the present invention may store multiple code images 131-133 for operating the various processors 105.

The master source 108 may distribute the code images 131-133 either by providing and broadcasting one code image, which may be all of the code images 131-133, for any code image request, or may optionally respond to a specific request for a specific code image. Thus, as an option, upon completion of the predetermined

DOCKET: TUC920000080-US1

time period, as determined by step 323, the master source 108, in step 330, determines whether the code image request received in step 320, and any code image request received in step 322, is a specific request for one of the multiple code images 131-133.

5 For example, an XY processor node 55, 255 may request a specific code image by accompanying the request with a specific identifier, such as an encoded number for the type of module.

If step 330 indicates that a specific code image has been requested, the master source selects that code image from
10 non-volatile memory 126 in step 331.

The master source 108 then, in step 335, distributes the code images 131-133 either by 1) providing and broadcasting one, possibly combined, code image for the code image request, or by 2) broadcasting only the requested one of the different code
15 images selected in step 331. A redundant master source may also request a code image 130, which is selected by master source 108 in step 331 and supplied in step 335.

The broadcast code images are supplied at the interface 112 of each of the processors 105 in step 336. In accordance with an
20 aspect of the present invention, if the processor is in the fully operational state and running, that operation will not be interrupted or disturbed, and, in step 337, the code image is ignored and not received. If the processor is in the minimally

DOCKET: TUC920000080-US1

operational state, the broadcast code image(s) are received in step 338. The code for making the determination of step 336 and conducting step 337 or step 338 may be part of the boot program of non-volatile memory 113, or part of the operational code of
5 RAM 114.

In an optional situation where, either the master processor 108 is capable of broadcasting only a single code image in step 331 at one time, the specific code image of a first specific request for a code image of step 311, may be broadcast even
10 though other processors requested different code images, or a single code image comprising a combination of code images 131-133 is broadcast, then, each of the processors 105 must determine whether it is receiving a correct code image. If that option is available, step 340 determines that the processor must select the
15 correct code image. The determination whether the code image is correct, or the correct code image is included in the broadcast, is made in step 343. If the code is correct, it is selected in step 345, and, if not, it is ignored in step 346 and the process cycles back to step 338 to receive the next broadcast.

20 Thus, where the one broadcast code image comprises a combination of different code images, the requesting processors 105 additionally select, in step 343, the code image that is correct for the processor. Where the broadcast code image is one

DOCKET: TUC920000080-US1

of the different code images 131-133 selected by the master source 108 in response to a specific request, each of the requesting processors 105 determines, in step 343, whether the broadcast code image is correct for the processor, and selects
5 the broadcast code image for implementation if the determination is that the code image is correct for the processor.

Whether selected by step 345, or if no selection is required from step 340, the code image is stored in RAM 114 and implemented in step 350. As is known to those of skill in the
10 art, the code image may overwrite any data previously in RAM 114. As an example, the boot program of non-volatile memory 113 may begin execution of the code image in step 350. The code for making the determinations of steps 340 and 343, the selection and ignoring steps 345 and 346, and the storage of the code image are
15 also part of the boot program of non-volatile memory 113.

In the example of an embedded system of FIGS. 3 and 4, the master source may comprise work processor 52, and the processors which may request code images may comprise the processors associated with the data storage drives 15, or may comprise the
20 communication processor nodes 50, 155, 250, the work processor node 252, the XY processor nodes 55, 255, or the operator panel processor nodes 59, 259. Also, a backup master source may be provided, for example, at the work processor node 252.

DOCKET: TUC920000080-US1

As the result, the present invention, without requiring a high capacity NVRAM, PROM or ROM at each processor for storing the entire code image during a power off situation, provides the correct code images to each of the requesting processors of a
5 multi-node processing network.

In embedded systems which, for example, require continuous operation, such as an automated data storage library, the modules are preferably redundant, or "hot-swappable", so that they may be repaired, replaced or upgraded as needed without requiring
10 turning off the entire system. In the example of FIGS. 3 and 4, the robot accessors may be redundant, such that work processor 252 contains the same code image as work processor 52, and may also comprise a backup master source, if needed. Also redundant are the XY processor nodes 55, 255. The operator panel processor
15 nodes 59, 259, and the communication processor nodes 50, 155, 250 may also be redundant, as may the set of data storage drives 15 in each of the frames 11-13.

In a non-redundant system, a code image for the processors may be updated by a reboot of the processors that are on the
20 network in accordance with the present invention, such that the processors will request the code image, and the master source will supply the updated code image.

Referring additionally to FIG. 6, in another aspect of the present invention, in which the processors are provided for

DOCKET: TUC920000080-US1

modules of a redundant system, such as an automated data storage library 10, having at least two sets of modules, the master source receives and stores the update code image(s) in step 360. A first set of the redundant modules is taken offline and
5 rebooted in step 362, while at least a second set continues in operation in step 363. As an example, work processor 52 may comprise the master source, and the modules of the high availability frame 13 may be taken offline and rebooted, comprising processors 252, 255, 250, 259, and the processors
10 associated with the data storage drives 15 in the frame. In step 363, the modules of the base frame 11 and extension frame 12 may continue operation, comprising processors 52, 55, 50, 155, 59, and the processors associated with the data storage drives 15 in the frames. In one embodiment, the master source 52 may send
15 messages to cause the processor to reboot and run the boot code as though the module is being reset after a power-on or is being reinitialized as discussed above, and, in another embodiment, the modules are rebooted by an operator or an external command.

In step 366, the boot code is executed, such that the
20 processors of the modules reach the minimally operational state and request the code image from the network. The master source 52, upon receiving a code image request in step 370, waits a predetermined time period, the predetermined time period allowing

DOCKET: TUC920000080-US1

any additional processor of the first set of modules to reach the minimally operational state and request code images. The master source 52, upon completion of the predetermined time period, in step 373, broadcasts the code image(s) on the network, such that
5 the processors requesting the code image(s) each receive the code image(s) in step 374, store the correct code image, and execute the code image and become fully operational.

Then, in step 380, a remaining second set of redundant modules is rebooted, as above. For example, the modules of the
10 base frame 11 and extension frame 12 may be taken offline and rebooted, comprising processors 55, 50, 155, 59, and the processors associated with the data storage drives 15 in the frames. Master source 52, or 252, if redundant, remains operational, as do the modules that became operational in step
15 374. Upon being rebooted, the processors 105 of the remaining set of redundant modules, in step 382, reach the minimally operational state and request the code image from the network. The master source 52, in step 385, upon receiving a code image request, waits a predetermined time period, the predetermined
20 time period allowing any additional processor of the second set to reach the minimally operational state and request a code image. In step 388, upon completion of the predetermined time period, the master source 52 broadcasts the code image(s) on the

DOCKET: TUC920000080-US1

network. In step 390, the second redundant set of processors requesting the code image(s) each receive the code image(s), store the correct code image, and execute the code image and become fully operational. Steps 380-390 may be repeated for 5 further sets of redundant modules.

In this manner, the code images for the entire redundant system have been updated without requiring that the entire system be interrupted.

As the result, the present invention eliminates the need to 10 maintain and update code levels on individual processors.

While the preferred embodiments of the present invention have been illustrated in detail, it should be apparent that modifications and adaptations to those embodiments may occur to one skilled in the art without departing from the scope of the 15 present invention as set forth in the following claims.

I claim: